

# Speech on the Mac

Although it may surprise many people, the Mac is quite talented when it comes to speech. Its abilities fall into two categories: taking commands from your voice (called “speech recognition”), and reading text aloud, using a synthesized voice (known in the biz as “text to speech”). This free Appendix covers both topics.

## Speech Recognition

The Apple marketing machine may have been working too hard when it called this feature “speech recognition”—the Mac OS feature called PlainTalk doesn’t take dictation, typing out what you say. (For that, you need a program like MacSpeech Dictate, [www.macspeech.com](http://www.macspeech.com)—or, better yet, Dragon NaturallySpeaking, an amazing dictation program for Windows that you can run on the Mac as described in Chapter 8.)

Instead, PlainTalk is what’s known as a *command-and-control* technology. It lets you open programs, trigger AppleScripts, choose menu commands, trigger keystrokes, and click dialog box buttons and tabs—just by speaking their names.

Few people use PlainTalk speech recognition. But if your Mac has a microphone, PlainTalk is worth at least a 15-minute test drive. It may become a part of your work routine forever.

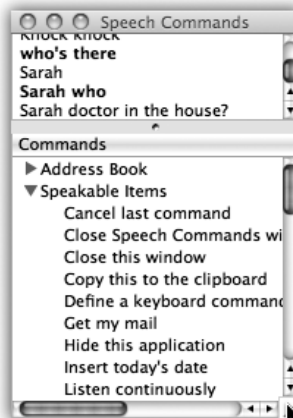
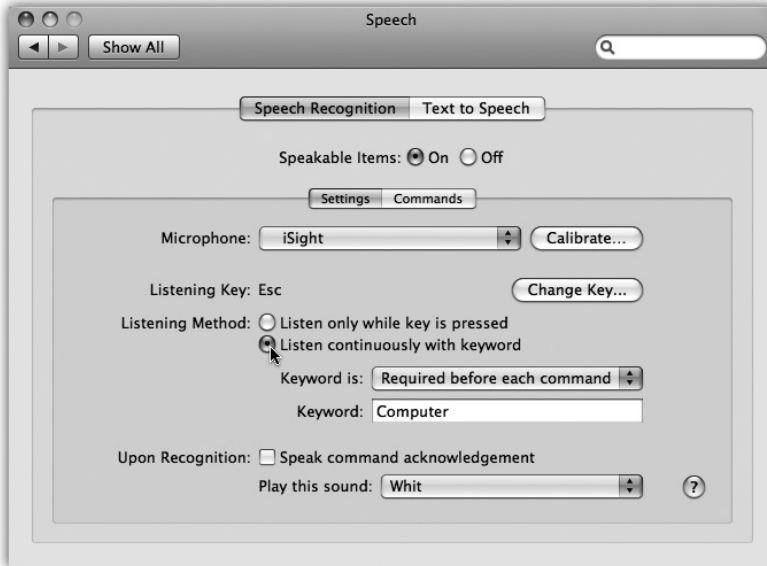
## Your First Conversation with the Mac

The on/off switch for speech recognition in Mac OS X is the Speech pane of System Preferences (Figure 1). Where you see “Speakable items” (on the Speech Recognition tab), click On.

## The Feedback window

Check out your screen: A small, microphone-like floating window now appears (Figure 2).

**Figure 1:** Turn listening on and off here. If you turn on “Listen only while key is pressed,” then the Mac pays attention to you only when you’re pressing a key (like Esc). As a convenience, it even pauses iTunes playback while you’re pressing the key. If you turn on “Listen continuously with keyword,” then you have to say a certain keyword to “get its attention” before speaking each command. In the Keyword blank, type the word you want the Mac to listen for as it monitors the sound from your mike.



**Figure 2:** Left: The Feedback window lacks the standard Close and Minimize buttons. If it's in your way, just double-click it (or say, “Minimize speech window”) to shrink it into your Dock. If you choose Speech Preferences from its bottom-edge triangle, you open the Speech Preferences window.

Right: Choosing Open Speech Commands window, of course, opens the list of things you can say.

The “Esc” in its center indicates the “listen” key—the key you’re supposed to hold down when you want the Mac to respond to your voice. (You wouldn’t want the Mac listening all the time—even when you said, for example, “Hey, it’s cold in here. Close the window.” Therefore, the Mac comes set to listen only when you’re pressing that key.)

You can specify a different key, if you wish, or eliminate the requirement to press a key altogether, as described in the next section.

When you start talking, you'll also see the Mac's interpretation of what you said written out in a yellow balloon just over the Feedback window.

### ***The Speakable Commands window***

The only commands PlainTalk understands are listed in the Speakable Commands window. (If it's not open, see Figure 2.) Keeping your eye on this window is essential, because it offers a complete list of everything your Mac understands. As you can see, some of the commands represent shortcuts that would take several steps if you had to perform them manually.

Here are a few examples of what you'll find in the list at first:

- **Phone for [Steve Jobs].** Displays Steve Jobs's phone number in huge digits across your screen—the fastest way yet to look up a number of somebody in your Address Book program.

You can also say “Chat with Steve Jobs” (or whomever) to begin a new chat session in iChat with that person, “Mail this to Steve Jobs” to send the current document as a Mail attachment, or “Meet with Steve Jobs” to add an iCal appointment with this person's information attached.

- **Set alarm for 30 minutes.** (You can say any number of minutes or hours—“Set alarm for 2 hours,” or whatever.) The Mac asks you to type a little message, which will serve as the iCal dialog box that appears at the specified time.

---

**Tip:** Check the complete list of iCal and Address Book commands in the Speech Commands window.

---

- **Close this window.** Closes the frontmost window instantly.
- **Empty the Trash.** Works only when you're in the Finder.
- **Switch to Safari.** Brings Safari to the front. (Actually, you can say, “Switch to” and then the name of *any* running or recently used program.)
- **Quit all applications.** Saves you the trouble of switching into each program and choosing Quit.
- **Open the Speech Commands window or Show me what to say.** Opens the Speech Commands window, of course.
- **What day is it?** Tells you the date.
- **Tell me a joke.** Begins a pathetic/funny knock-knock joke. You've got to play along, providing the “Who's there?” and “So-and-so *who?*” answers.

Mac OS X *updates* the listing in the Speech Commands window in real time, according to the context. When you switch from one program to another, you see a list of the local commands that work in the new program. You'll discover that when you use the “Tell me a joke” command, for example, you don't necessarily have to say, “Who's

there?” You can also say, “Stop,” “Go away,” or “Stop with the jokes!” (It must really be fun to work at Apple.)

### ***Speaking to the Mac***

Finish up on the Speech pane of System Preferences. Use the Microphone pop-up menu to specify which microphone you’ll be using (if you have a choice). Click Calibrate to adjust its sensitivity.

Now you’re ready to begin. While pressing the Esc key (if that’s still the one identified in the Feedback window), begin speaking. Speak normally; don’t exaggerate or shout. Try one of the commands in the Speakable Commands list—perhaps “What time is it?” If your mike is set up properly, the round Feedback window displays animated sound waves when you speak.

## **Customizing Speech Recognition**

You can tailor the speech recognition feature in two ways: by adjusting the way it looks and operates, and by adding new commands to its vocabulary.

### ***Changing when the Mac listens***

Having the microphone “open,” listening full-time, is an invitation for disaster. Everyday phone conversations, office chatter, and throat clearings would completely bewilder the software, triggering random commands.

Therefore, you must explicitly *tell* the Mac when you’re addressing it. When you first turn on the speech recognition feature, the Mac expects you to alert it by pressing a key, like Esc, when you speak.

---

**Tip:** To change the key you hold down when you want the Mac to listen, visit the Speech pane of System Preferences; click the Speech Recognition tab; click Settings; and then click Change Key. A little message prompts you to press the keyboard key you’d prefer to use. Your choices are Esc, Delete, F5 through F12, or the keys on your numeric keypad—with or without the Shift, Control, or Option keys.

---

If you’d rather not have to press some key whenever you want the computer’s attention, click the other option in this pane, “Listen continuously with keyword.” Now to get the computer’s attention, you must speak the keyword—which you type into the Keyword box—before each command. For example, you might say, “Computer, open AppleWorks,” or “Hal, what day is it?”

The word you specify appears in the middle of the round Feedback window.

---

**Note:** This method of getting the computer’s attention is less reliable than the push-a-key-to-talk system. Especially if you name the computer “Hal.” Although that’s hilarious in theory, polysyllabic words work better in practice.

---

By using the “Keyword is” pop-up menu, meanwhile, you can specify how big your window of opportunity is:

- **Optional before commands.** If you work alone in a quiet room, this is the choice for you. It means you don't have to press a key *or* say the Mac's name when issuing a voice command. *Everything you say* is considered a command.
- **Required before each command.** Nothing you say is interpreted as a command unless you say the computer's keyword first, as in, "Macintosh, switch to Microsoft Word."
- **Required 15 seconds after last command, Required 30 seconds after last command.** Sometimes you want to issue several commands in a row and would feel foolish saying, "Computer, close all windows. Computer, empty the Trash. Computer, switch to Calculator." When you turn on this option, you can say the keyword just once; all commands you issue in the next 15 or 30 seconds "belong to" that first salutation. The push-to-talk key and the spoken keyword, in this case, serve as a master on/off switch for the Mac's listening mode.

---

**Tip:** If you're not using the push-to-talk method, you can still turn speech recognition off temporarily by saying, "Turn on push to talk." (Now the Mac listens to you only when you're pressing the designated key.) When you want to return to listening-all-the-time mode, say, "Listen continuously."

---

### **Changing the feedback**

Another set of options on the Speech Recognition tab governs what the Mac does when it understands something you've said. For example:

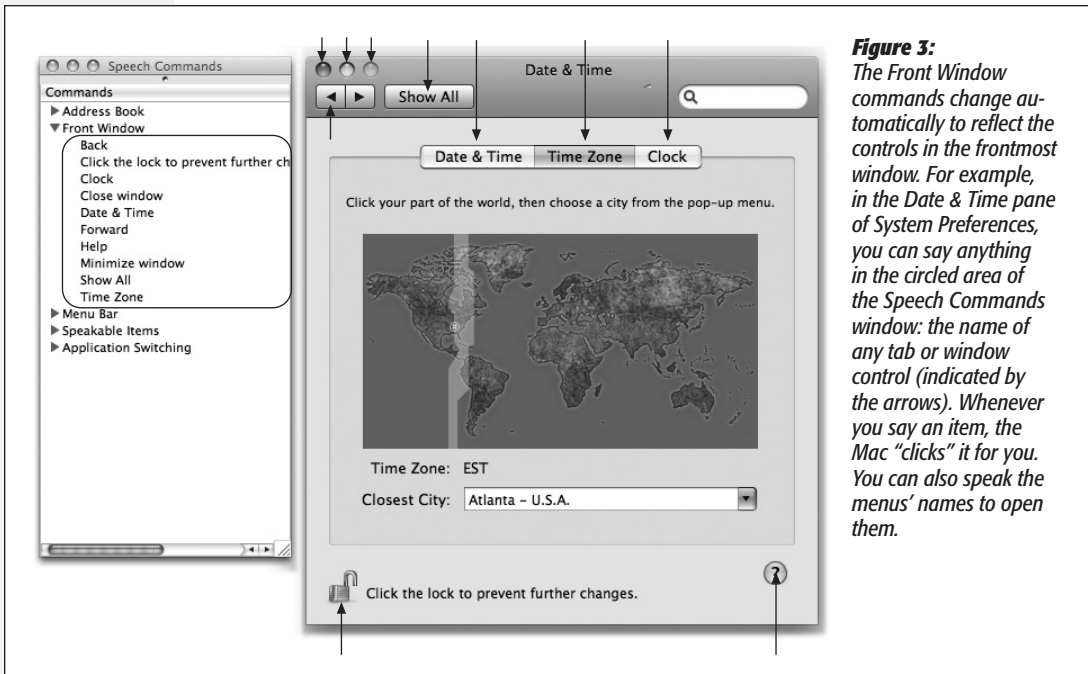
- **Play sound.** The Mac generally makes a sound whenever it recognizes something you've said. Use this pop-up menu to control which of your built-in beeps you want it to use—or choose None.
- **Speak command acknowledgement.** Sometimes the Speech Feedback window shows you a message of its own. When you use the "Empty the Trash" command, for example, text in the Feedback window may inform you that a locked item prevents the emptying. The Mac generally reads this text aloud to you; turn this checkbox off if you'd rather have the Mac be silent.

### **Triggering menus by voice**

On the Speech pane of System Preferences, click the Speech Recognition tab, and then click the Commands minitab. Here you find a list of the command categories that Speakable Items can understand. As you turn each checkbox on or off, watch the Speech Commands window. Giant swaths of commands appear or disappear as you fool with these checkboxes, giving you a good indication as to their function. Here's a rundown:

- **Address Book.** These commands let you look up numbers; add appointments to iCal; set up alarm reminders for yourself; mail things to people; and begin text, audio, or video chats with people whose names are already in your Address Book.
- **Global Speakable Items.** This is the master list of Speakable Items, shown in Figure 2.

- **Application-Specific Items.** Certain Mac OS X programs come with preset lists of commands that work only when you're in the relevant program. For example, whenever you're in the Finder, you can say, "Empty the Trash," "Go to my Home directory," "Hide the Dock," "Minimize all windows," "Make a new folder," and so on. When this checkbox is off, the Mac no longer recognizes any of these handy commands.
- **Application Switching.** This is the command category at the bottom half of the Speech Commands list—"Switch to Address Book," "Switch to AOL," and so on.
- **Front Window.** In your Speech Commands window, note the appearance of a new category of commands, called Front Window. The idea here is to provide you with quick speech-recognition access to the most prominent buttons, tabs, and icons in whichever window is before you. Figure 3 elaborates on the idea.



**Figure 3:** The Front Window commands change automatically to reflect the controls in the frontmost window. For example, in the Date & Time pane of System Preferences, you can say anything in the circled area of the Speech Commands window: the name of any tab or window control (indicated by the arrows). Whenever you say an item, the Mac “clicks” it for you. You can also speak the menus’ names to open them.

- **Menu Bar.** This command lets you open menus (in the menu bar) by speaking their names.

Once you say its name (“File menu,” for example), the menu opens. Now you can say any command in the open menu (“New Playlist,” “Save,” or whatever). The Menu Bar category of the Speech Commands window changes to remind you of what you can say at any given moment.

The combination of Front Window and Menu Bar commands lets you do quite a bit of work on your Mac without ever needing the mouse or keyboard.

### ***Improving the PlainTalk vocabulary***

By putting an alias of a favorite document, folder, disk, or program into your Home→Library→Speakable Items folder, you can teach PlainTalk to recognize its name and open it for you on command. You can name these icons anything you want.

You can also rename the starter set that Apple provides. You'll have the best luck with multiword or polysyllabic names—"Microsoft Word," not just "Word."

One kind of icon PlainTalk can open is an *AppleScript* icon. If you open your Home folder→Library→Speech→Speakable Items, you'll discover that most of the built-in Speakable Items icons are, in fact, AppleScript icons. The point is you can make PlainTalk do almost anything you want, especially in the Finder, simply by creating AppleScripts and putting them into the Speakable Items folder.

### ***Application-specific commands***

Most of the preinstalled PlainTalk commands work in any program. You can say, for example, "Open iTunes" to launch iTunes from within any program.

However, you can also create commands that work only in a specific program. They sit in your Speakable Items→Application Speakable Items folder, inside individual application-name folders. For example, Mac OS X comes with commands for Safari that include Go Back, Go Forward, and Page Down.

If you get good at AppleScript, you can create your own application-command folders in the Speakable Items→Application Speakable Items folder.

Then open the program for which you want to create special commands and say, "Make this application speakable." The Mac creates a folder for the program in the Speakable Items folder; fill it with the AppleScripts you've created. (Not all programs can be AppleScript-controlled.)

---

**Note:** If you give an application-specific icon the same name as one of the global commands, the Mac executes the application-specific one—if that program is running.

---

### ***PlainTalk tips, tricks, and troubleshooting***

When you're creating new commands, click the Helpful Tips button at the lower-right corner of the Commands pane (of the Speech Recognition pane of the Speech pane of System Preferences).

## **The Mac Reads to You**

So far in this chapter, you've read about the Mac's listening ability. But the conversation doesn't have to be one-way; it's even easier to make the Mac *talk*.

In Snow Leopard, most Apple programs have a Start Speaking command built right in: Safari, Mail, TextEdit, iChat, Stickies, Pages, and so on. Just Control-click (or right-click) inside a window full of text and, from the pop-up menu, choose Speech→Start Speaking. How cool is that? Your Mac can read your email or a Web article to you while you're getting dressed.

You can add a Speak command in FileMaker Pro scripts. Mac OS X's Chess and Calculator programs can talk back, too.

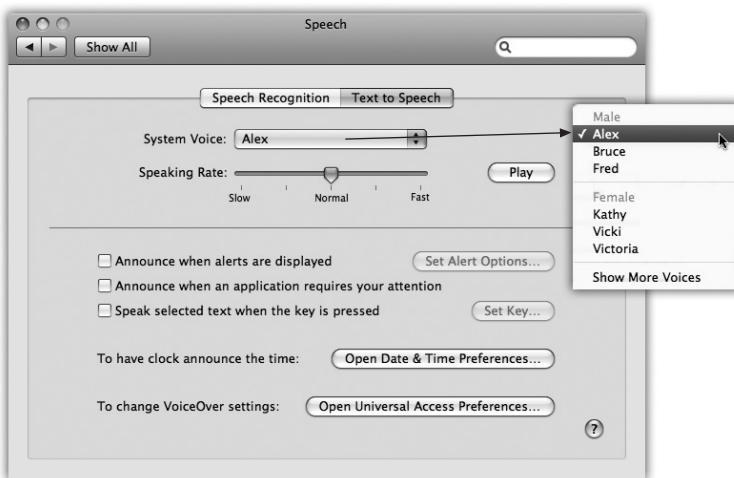
But that's kid stuff. Truth is, the Mac can read almost anything you like: text you pass your cursor over, alert messages, menus, and *any text document in any program*. It can speak in your choice of 24 synthesizer voices, ages 8 to 50. Most read with a twangy, charmingly Norwegian accent—all but Alex, Apple's newest voice, which sounds scarily like a professional human voice-over artist.

## Setting Up the Mac's Voice

To configure the way the Mac talks, revisit the Speech pane of System Preferences. Click the Text to Speech tab at the top of the window. As you can see in Figure 4, you can control which of the Mac's voices you want your computer to use, as well as how fast it should speak.

**Tip:** Five of the voices sing rather than speak. Good News sings to the tune of *Pomp and Circumstance*, otherwise known as the Graduation March. Bad News sings to the tune of Chopin's *Prelude in C Minor*, better known as the Funeral March. Cellos sings to the tune of Grieg's *Peer Gynt* Suite. Pipe Organ sings to the tune of the Alfred Hitchcock TV theme. Bells sings the typical church-bell carillon melody.

In other words, these voices sing whatever words you type to those melodies. (To hear the melody in its entirety, don't use any punctuation.)



**Figure 4:**

At the outset, you see only six voices—the ones that sound the most human. Choose Show More Voices to see the complete list of 24.

Then, for 15 minutes of hilarity, try clicking the voices in turn to hear sample sentences. Drag the slider to affect how fast each one speaks. (Clearly, Apple's programmers had some fun with this assignment.)

Here are all the different occasions when the Mac can talk to you.

### ***Announce when alerts are displayed***

If you turn on this checkbox, you can make the Mac read aloud error messages and alert messages that may appear on your screen. If you click the Set Alert Options button, you find these useful controls:

- **Voice.** Use this pop-up menu to specify the voice you want reading your error messages. (It doesn't have to be the same as the standard Mac default voice that's used for other purposes.)
- **Phrase.** Use this pop-up menu to specify which utterance the Mac speaks before the actual error message—for example, “*Excuse me!* The Trash could not be emptied,” or “*Attention!* The document could not be printed.”

If you choose “Next in the phrase list” or “Random from the phrase list” from this pop-up menu, you'll never hear the same expletive twice. Better yet, choose Edit Phrase List to open a dialog box where you can specify your own words of frustration. (Apple Computer, Inc., is not liable for any trouble you may get into with people in neighboring cubicles.)

- **Delay.** The ostensible purpose of the Talking Alerts feature is to get your attention if you've wandered away from your Mac—mentally or physically. The chances are slim, but an urgent problem might occur that, if left undetected, could land you in trouble. (A 500-page printout brought to its knees by a paper jam comes to mind.)

In other words, if you're still sitting in front of your Mac, you may not need the Mac to speak to get your attention; you could simply read the onscreen message. That's why you can set this slider to make the Mac wait, after the error message appears, for up to a minute before trying to flag you with its voice. That way it won't harangue you unnecessarily.

### ***Announce when an application requires your attention***

A jumping Dock icon means that the program is trying to get your attention. It might be because your printer is out of paper, or it might be because you've tried to shut down but one program still has open, unsaved documents. In any case, now the Mac can tell you, in so many words, which program needs some loving care.

### ***Speak selected text when the key is pressed***

As promised, this is the feature that lets you hear any Web page read to you, any email message, any sticky note—a welcome break for sore eyes. Furthermore, if you do any kind of writing at all, you'll discover the value of having your material read back to you out loud; it's a proofreading technique that reveals all kinds of typos and wordos that you wouldn't catch just by reading, even if you read it 50 times.

The very first time you turn on this checkbox, the Mac prompts you to specify a keystroke. (Later, you can always change the keystroke by clicking Set Key.) Choose a keystroke that doesn't conflict with the program you're using, like Control-T.

Now go to the program where you'd like the reading to happen. Highlight some text (or press **⌘-A** to select all of it). Then press the keystroke you specified. The Mac begins reading it aloud immediately. To interrupt the playback, press the same keystroke again.

And be glad you were alive to see the day.

---

**Tip:** If you have an iPod or iPhone, your Mac can convert any text or word processing document into a spoken recording so you can listen to it when you're on the go. To pull this off, see page 265.

---

## VoiceOver

The Mac has always been able to read stuff on the screen out loud. But Apple has taken this feature light-years farther, turning it into a full-blown *screen reader* for the benefit of people who can't see. VoiceOver doesn't just read every scrap of text it finds on the screen—it also lets you control everything on the screen (menus, buttons, and so on) without ever needing the mouse. And in Snow Leopard, VoiceOver has taken another *enormous* leap forward.

Here's a little bit of what's new:

- **Gestures.** On a Mac with a multitouch trackpad, you can interact with what's visible on the screen by using the trackpad itself as a map of the current window or screen area; VoiceOver speaks whatever's under your finger. You can flick in any direction to move to the next thing on the screen. You hear a sound whenever VoiceOver finds blank spaces on the screen, all in the name of helping you “feel” where everything is.
- **Braille monitors.** Yes, there are actually Braille “displays”—essentially flat, touchable panels with dots that rise through holes in a flat surface as you move the Mac's cursor, permitting you to read what's on the screen. Snow Leopard works with more of these contraptions, including wireless ones, and can broadcast to a classroom full of them simultaneously.
- **VoiceOver Web browsing.** Snow Leopard can read entire Web pages, navigate Web tables, hop from link to link, and so on. You can navigate by pressing arrow keys, using the mouse, or using the trackpad.
- **More flexibility.** Snow Leopard brings more customizability, more settings, and more help getting started, thanks to an interactive tutorial.

Obviously, learning VoiceOver is a *huge* task that can take days or weeks—but if it's your ticket to being able to use a computer at all, you'll probably be happy to have such a full-fledged monster of a program.

A full guide to VoiceOver could easily fill 40 pages—and does. Download the free PDF appendix to this chapter, “VoiceOver: The Missing Manual” from this book's “Missing CD” page at [www.missingmanuals.com](http://www.missingmanuals.com).